

Workshops and Conferences

The ISCR-sponsored or co-sponsored 10 scientific workshops in FY 2004. Two of these were hosted locally and exclusively by the Laboratory. The rest were hosted in cooperation with other organizations, such as the Society for Industrial and Applied Mathematics (SIAM), the Institute for Pure and Applied Mathematics (IPAM), the Department of Homeland Security, or Argonne National Laboratory and held off site. Some ISCR workshops are one-of-a-kind exploratory workshops that assemble experts to scope out possible new programs. Others have become part of the fabric of their disciplines and are held at regular intervals. In each case, there is a vital LLNL interest and typically, several Laboratory researchers participate.



The 2003 International Conference on Preconditioning Techniques for Large Sparse Matrix Problems in Scientific and Industrial Applications

Dates October 27–29, 2003

Location Napa, California

The 2003 International Conference on Preconditioning Techniques for Large Sparse Matrix Problems in Scientific and Industrial Applications (Preconditioning 2003) was the third conference of its kind to focus on preconditioning techniques for solving sparse matrix problems. The first conference took place in Minneapolis, MN, in June 1999 and the second in Tahoe City, CA, at the end of April 2001. The first conference drew close to 100 participants and the second about 70 participants, while the third drew about 80 participants.

One of the characteristic themes of the meeting is its emphasis on real-life (“industrial”) problems. For this reason, all three conferences enjoyed a healthy balance between academia and industry/government labs in its mix of participants. In addition, there is a rather important contingent of participants from overseas (mainly Europe). The preconditioning meetings have been quite successful and are now being viewed by the community as the premier specialized conference on preconditioners.

The Napa conference featured 7 invited speakers, 29 contributed papers and 8 posters. One of the goals of the meeting is to foster dialogue

between practitioners and academics. As a rule, the organizing committee gives a charge to the program committee to nominate invited speakers with a goal of keeping a good balance between the number of talks on “methods” and those on “applications.” The organizing committee then finalizes the selection to reach this goal. This particular meeting reached a good balance. There were four invited talks on applications and three invited presentations on algorithmic aspects.

One feature that distinguished this meeting from the previous two was that we allowed more parallel sessions. Previously, few (Tahoe City) or zero (Minneapolis) parallel sessions were scheduled. This change was necessary to increase the number of presentations as there were many high-quality abstracts.

There was a good mix of attendees from academia, research laboratories, and industries. In particular, the DOE was well represented, with more than 20 participants from all of the major DOE laboratories (ANL, LBNL, LLNL, LANL, and SNL). Industry participation was also relatively strong. Overseas participants came from Belgium, China, Denmark, France, Germany, Great Britain, Israel, Italy, Japan, The Netherlands, Russia, and Tanzania.

Multi-Algorithm Methods for Multiscale Simulations Workshop

Dates

January 14–16, 2004

Location

Livermore, California

The Multi-algorithm Methods for Multiscale Simulations Workshop was held January 14–16, 2004, at the Hilton Garden Inn in Livermore, California. The event was hosted by CASC, ISCR and LLNL. It was sponsored by the Institute for Terascale Simulation (ITS) on behalf of the ASC Program. In all, 45 attendees, including 26 from U. S. Department of Energy laboratories and the balance from academia and industry, participated in the two-and-one-half day workshop.

Multiscale simulation is a central emerging numerical modeling paradigm for many science and engineering problem areas at LLNL and throughout the scientific community. Application areas include: materials design (nano-wires in computer chips, photonics, micro-electromechanical systems); biological systems (protein docking); medicine (drug delivery systems); and nuclear and aerospace technologies where materials failure and response in severe environments is a primary concern (dislocation patterns in fatigue and creep, surface roughening and crack nucleation in fatigue, etc.). Important scales in such problems range from macroscopic, where continuum models based on differential equations are usually employed, to atomistic, with quantum mechanical models applied at the finest resolution, plus all intermediate scales.

The Multi-algorithm Methods for Multiscale Simulations Workshop comprised five half-day sessions on various multiscale topics. Sessions were sorted by the organizers into methods and applications in solids, liquids, and gases and cross-cutting techniques. Each session concluded with a panel discussion in which the speakers and audience interacted richly and informally.

An important aim of the workshop organizers was to gather practitioners of multi-algorithm methods from various fields to find common ground in their work and engage in intellectual cross-fertilization. For example, materials scientists and gas dynamicists are not accustomed to working

closely together, let alone with the mathematicians and computer scientists also present, since traditional conferences and meetings are usually focused on specific disciplines or problem areas. Many of the workshop participants explicitly stated that they found the workshop format and technical exchanges refreshing and informative. Several participants expressed strong interest in making the workshop an annual event.

Common challenges identified at the workshop included locating relevant scale boundaries and in grafting together representations of the solution from different methods in such boundary regions in which both are valid, in order to model a global system for which no single method is everywhere valid or efficient. Participants also discussed important challenges in understanding physical and mathematical error analysis in hybrid computational models and implementation challenges associated with complex models for large-scale parallel platforms.

An especially important outcome of the workshop is that DOE and academic researchers gained familiarity with many commonalities, as well as differences, among the problems on which they work. For example, it was revealed that understanding dissipative processes in solid mechanics (e.g., dislocations) can benefit from work done in modeling liquids. Also, the principle of entropy production used to understand the breakdown of continuum models in liquid and gas simulations appears to possess similarities with the principle of power dissipation used in solid mechanics. Other important areas identified as requiring further exploration included methods for bridging widely disparate time scales in hybrid simulations and methods for fluid-structure/gas-surface interactions, as well as incorporating more detailed chemistry in biological and nanoscale problems.

https://www.llnl.gov/casc/workshops/multiscale_simulations

Copper Mountain Conference on Iterative Methods

Dates

March 28 – April 2, 2004

Location

Copper Mountain, Colorado

The Copper Mountain Conference Series, held annually in early April at Copper Mountain, CO, alternates subjects between Multigrid Methods in odd-numbered years and Iterative Methods in even-numbered years. It represents an important forum for the exchange of ideas in these two closely related fields.

The Copper Mountain Conference on Iterative Methods was held March 28 – April 2, 2004. A total of 185 mathematicians from all over the world attended the meeting, which began with a reception on Sunday, March 28. During the following 5 days of the meeting, 131 talks on current research topics were given. Talks were organized into the following sessions.

1. Multigrid Solvers and Algebraic Multigrid
2. Saddle-Point Solvers
3. PDE Methods
4. Preconditioning Methods
5. Eigenvalue Methods
6. Multi-Physics Solution Methods
7. Krylov Subspace Methods
8. First Order System Least Squares Methods (FOSLS)
9. Nonlinear Solvers
10. Continuation Methods
11. Stochastic Systems
12. Parallel Algorithms
13. Applications
14. Software

In addition to the regular sessions, three evening workshops were offered. Monday night's workshop was a mini-symposium organized by Henry Tufo, who represented NCAR and the University of Colorado. Tuesday night, Michael Heroux from SNL organized a workshop on Sandia's Trilinos project. Wednesday night, a workshop organized by Eldad Haber from Emory University highlighted PDE-Constrained Optimization.

The sessions were all very well attended. The Copper Mountain Conference Series is known for having a very relaxed atmosphere and for fostering open, active discussions. This collaborative environment has characterized the meeting since the series began in 1983, and is one of the reasons many attendees come back repeatedly.

A student paper competition was held to stimulate student participation in the Conference. A panel of judges made up of members of the Program Committee selected the winners: Yair Koren (Technion, Israel); Ruth Holland (Oxford University, England); and Andrei Draganescu (University of Chicago).

<http://amath.colorado.edu/faculty/copper/2004>

Statistics and Practical Applications of Data Mining: Highlights from SDM04

Dates

April 22–24, 2004

Location

Orlando, Florida

The Fourth SIAM International Conference on Data Mining, held in Orlando, FL, April 22–24, 2004, continued the tradition of providing an open forum for the presentation and discussion of innovative algorithms, as well as novel applications of data mining. A record number of paper submissions this year marked not only a growing interest in the field, but also a greater acceptance of the conference among data mining researchers and practitioners.

Student authors accounted for a large percentage of the accepted papers, and their papers were reviewed under the same stringent guidelines as regular papers. The best student paper award was given to Martin Law from Michigan State University for his work on manifold learning. The award for the best algorithms paper went to a team from the University of Texas at Austin for their work on clustering, while the best applications paper was on enhancing communities of interest by a team from AT&T Laboratories.

A running theme of the conference was the practical application of data mining, including opportunities in various problem domains and practical lessons learned by those solving real data analysis problems in these domains. This was reflected in the topics covered in the three tutorials: analysis of patients' medical data, data mining for computer security, and mistakes commonly made in data mining and ways to avoid them.

In an industry–government session, speakers discussed problems encountered in the telecommunications industry, the role of information visualization, and data mining in such diverse domains as aviation safety and security, performance of computer networks, and earth sciences. Applications of data mining were also the subject of three of the keynote talks: Sara Graves of the University of Alabama at Huntsville considered issues of data usability; David Page of the University of Wisconsin Medical School elaborated on data mining questions

raised by biology data; Ted Senator from the Defense Advanced Research Projects Agency (DARPA) discussed “connecting the dots.” The increasing importance of homeland security was also reflected in many of the conference workshop topics, which ranged from link analysis, counterterrorism and privacy, to data mining in resource-constrained environments. More traditional topics, such as bio-informatics, mining of scientific and engineering data sets, and high-performance and distributed mining, also continued to attract participants.

Conference attendees clearly welcomed the focus on applications, which led to animated discussions in the industry-government presentations. One workshop speaker took the tutorial by John Elder on common mistakes in data mining to heart; she did some real-time editing of her presentation to point out the mistakes in her application domain, such as a lack of caution in sampling the data and discounting pesky cases though they might reveal a larger problem in the data.

A new aspect of this year's conference was the increasingly important role of statistics in data mining. Keynote speaker Chris Bishop of Microsoft Research: Cambridge discussed recent advances in Bayesian inference techniques and several technical sessions focused on statistical techniques in data mining. This connection between statistics and data mining will be exploited further in the next conference in the series (scheduled for Newport Beach, April 21–23, 2005), which will be co-sponsored by the American Statistical Association and SIAM (<http://www.siam.org/meetings/sdm05/>). We encourage statisticians and data miners to submit papers and attend the conference, and help us narrow the gap between the two fields to bring together the best of both worlds.

The proceedings of SDM04, including the keynotes and the presentations at the industry/government session, are available on-line at <http://www.siam.org/meetings/sdm04>.

Department of Homeland Security Advanced Scientific Computing Workshops

Dates

May 12, 2004

September 22–23, 2004

Locations

Washington, DC

Alexandria, Virginia

As part of the Department of Homeland Security (DHS) Advanced Scientific Computing program, ISCR co-hosted three requirements-gathering workshops during FY 2003. The first of these, the Advanced Scientific Computing Requirements Workshop was held October 8-9, 2003; a summary of this workshop can be found in the *ISCR 2003 Annual Report*. The two additional DHS workshops held this year were the Incident Management Simulation Workshop and the Data Sciences Workshop. The Krell Institute participated in the development of the content of these workshops and handled all workshop logistics and developed the workshop Web sites.

The Incident Management Simulation Workshop was held on May 12, 2004 at the Westin Grand Hotel in Washington, DC. This workshop brought together senior representatives of the emergency response and incident management communities with modeling and simulation technologists from DOE laboratories. The workshop provided an opportunity for incident responders to describe the nature and substance of the primary personnel roles in an incident response, identify current and anticipated roles of modeling and simulation in support of incident response, and begin a dialog between the incident response and simulation technology communities that will guide and inform planned modeling and simulation development for incident response.

The workshop was a joint effort of the Advanced Scientific Computing Program and the DHS Emergency Preparedness and Response

Portfolio, both elements of the DHS Science and Technology Directorate. Based on the interactions at the workshop, a panel of computational science technologists prepared a summary report on incident management practice and the potential roles that computational simulation might play in supporting incident management. In particular, the panel prepared a summary of simulation capabilities that are relevant to incident management training and recommendations for the use of simulation in both incident management and in incident management training. In addition, the final report discusses areas where further research and development will be required to support future needs in this area.

The DHS Data Sciences Workshop was held September 22–23, 2004 at the Hilton Old Town in Alexandria, VA. The purpose of this workshop was to thoroughly review the data sciences mission needs of DHS and to identify specific mathematics and computer science research and development (R&D) topic areas required to address those needs. During the workshop, approximately 50 invited participants representing DHS, DOE and its national laboratories, academia and industry, identified specific R&D topic areas in the data sciences, their ties to the mission needs of DHS, and the potential impact of the proposed R&D. This effort will specify five years of relevant research topics in the data sciences area to support the Threat and Vulnerability Testing and Assessment Portfolio of DHS. These activities are important since they will immediately feed into the current DHS planning process for FY 2006.

Domain-Specific Languages for Numerical Optimization

Dates

August 18–20, 2004

Location

Argonne, Illinois

On August 18–20, 2004, ANL hosted a workshop on Domain-specific Languages for Numerical Optimization, co-sponsored by LLNL. There were 36 participants, including students, faculty, and staff from 12 universities, plus scientific staff from ANL, LLNL, and Sandia. The purpose of the meeting was to bring together experts in programming languages and compilers together with experts in numerical optimization and partial differential equations (PDEs) to stimulate discussion on the design and implementation of next-generation domain-specific languages for numerical optimization, with an emphasis on stochastic optimization and PDE-constrained optimization.

The participants discussed the design and implementation of current languages for numerical optimization, called modeling languages. These languages are mostly declarative, but take on an imperative flavor when a function or its derivatives must be evaluated. Native data types include scalars, sets, and ordered sets. Typically, the models and data are “compiled” into an internal representation or bytecode that is then interpreted. Several examples were given in which a problem that might have taken thousands of lines of code to express in Fortran or C required only 30–100 lines of code in AMPL or GAMS.

On the other hand, two examples were cited where converting a problem from a modeling language to C or Fortran (plus an FFT library in one case) resulted in a hundred- or thousand-fold speedup. Although the granularity of objects in modeling languages is typically much finer than that in other domain specific languages, it seems likely that static or dynamic compilation could provide the performance needed for some large problems without sacrificing the expressiveness of the modeling languages.

One of the obstacles to successfully extending modeling languages to support stochastic optimization is the wide variety of ways that randomness can enter an optimization problem and hence the many kinds of stochastic optimization problems. Even when the scope is restricted to a particular kind of stochastic optimization, multistage

linear recourse problem, the specification of a problem is nontrivial. Part of the challenge arises from the fact that multistage problems can lead to enormous scenario trees. However, for problems in which stochasticity enters in a structured manner, extending a modeling language with time (stage) information and mechanisms to specify the probability distributions for random variables may suffice. The primary obstacle to extending modeling languages to support PDE-constrained optimization is that effective mechanisms for specifying partial differential equations themselves have not been developed. However, emerging systems, such as Sundance, FIAT, and PETSc 3, offer some hope that effective mechanisms for specifying and solving PDEs can be developed.

Several systems for analysis and transformation of general-purpose languages, domain-specific languages, and meta-languages were presented. The DMS Software Reengineering Toolkit supports automated source code analysis and modification. It utilizes Unicode lexers, GLR parsers for arbitrary context free grammars, analysis via multipass attribute grammars, and conditional source-to-source transformations. The extensible C (xtc) system uses a packrat parser to support arbitrary syntactic extensions to C, AST transformation rules to reduce and optimize, and typing rules to support safety constraints. Several researchers presented their work in the area of telescoping languages and related techniques. These methods exploit domain-specific analysis and optimization to improve the performance of general purpose languages extended with domain-specific libraries. Examples of systems supporting this paradigm are ROSE, Broadway, and libGen. In many cases, these systems are able to achieve performance superior to voluminous hand-developed Fortran or C implementations using concise implementations in C++ or MATLAB. It was also demonstrated that generic programming techniques can also provide high performance and high levels of expressiveness.

<http://www-unix.mcs.anl.gov/workshops/DSLOpt/>

Short-Pulse Laser Workshop

Dates

August 25 – 27, 2004

Location

Pleasanton, California

For three days at the end of August 2004, 55 plasma scientists met at the Four Points by Sheraton in Pleasanton to discuss some of the critical issues associated with the computational aspects of the interaction of short-pulse, high-intensity lasers with matter. The workshop was organized around the following areas of key interest to the Laboratory.

- Laser propagation / interaction through various density plasmas: micro scale.
- Anomalous electron transport effects: from micro to meso scale.
- Electron transport through plasmas: from meso to macro scale.
- Ion beam generation, transport, and focusing.
- "Atomic-scale" electron and proton stopping powers.
- $K\alpha$ diagnostics.

Each area had a coordinator who drew up a list of questions, moderated discussions, and wrote a working group summary.

Many important problems in fast ignition are related to laser-plasma interactions, including laser propagation in the underdense corona plasma, laser hole-boring in the overdense plasma, laser absorption and energetic electron production at the critical surface, and electron transport in the mildly-dense plasma region. Participants worked out a set of benchmark computational simulation problems to compare their modeling capability in these areas.

Meso to macro scale electron transport discussions focused primarily on the correct method of initiating the electron beam. Several phenomenological techniques were discussed, such as injection at a plane in free space, promotion of ambient electrons, and the use of a ponderomotive force. It was generally agreed that the boundary conditions in the laser-plasma interaction (LPI) region were critical to the problem setup. A two-region approach, in which the LPI is simulated in the blow-off plasma and hybrid methods are used in the solid density material, might be a reasonable

intermediate step. Since most codes do not have a laser-deposition package, it is necessary to choose and standardize the beam parameters, and a set was proposed.

The ion beam generation, transport, and focusing group discussed the following questions, from general to application specific.

- What are the proton generation mechanisms?
- What are their efficiencies?
- How sensitive to resolution are the answers?
- What codes can be used?
- How does electron flow affect proton generation?
- How can we control the generation and focusing of the protons?
- What is the optimum proton energy for radiography?
- What are the qualities that set ions using these mechanisms apart from "standard" ion beams?
- What governs ion flux?
- What is the optimal distance of the "proton lens" from the target?

The "atomic-scale" stopping powers session discussed first the stopping power of relativistic electron beams (REB) with energies of 1 to 10 MeV, stopping in pre-compressed deuterium-tritium (DT) targets, and the stopping power of non-relativistic (NR) protons with energies of 1 to 100 MeV. They set benchmark problems for multiple scattering of REB on target ions and of multiple scattering of NR protons in thin foils of high-Z materials disposed in front of laser proton sources (LPS). Finally, they considered REB and NR proton stopping in strongly magnetized fast ignition targets.

The purpose of the $K\alpha$ diagnostic sessions was to discuss some of the progress being made in modeling $K\alpha$ emission in short-pulse petawatt laser experiments and to discuss with experimentalists some of their latest results.

The workshop was made possible by the joint financial support of the Institute for Laser Science and Applications and the ISCR at LLNL.

Computational Methods in Transport

Dates

September 11–16, 2004

Location

Tahoe City, California

The Computational Methods in Transport Workshop was devoted to providing a forum where computational transport researchers could discuss their methods, successes and failures across disciplinary boundaries. Typically, the numerical methods used in a given field are communicated to other researchers in that field. Rarely, however, are those methods communicated between one application domain of radiation transport and another.

For example, nuclear engineers and astrophysicists rarely attend the same meetings or read the same literature. The goal of the Computational Methods in Transport Workshop was to address this discrepancy and open channels of communication and cooperation so that (1) existing methods used in one field could be applied to other fields and (2) greater scientific resources could be mobilized to help solve outstanding problems.

Beginning on the afternoon of September 11, 2004 and ending with lunch on September 16, 2004, the workshop was held at the Granlibakken Conference Center in Tahoe City, CA. The first day of the meeting consisted of a series of one-hour talks reviewing one of the major fields represented at the workshop. The areas covered included astrophysics, atmospheric physics, mathematics, plant canopies, nuclear engineering, oceanography and high-energy density physics. One afternoon was reserved for a poster session where 30 posters were presented in a very lively and well-attended event. The following days were filled with focused 45-minute talks by each of the representative fields that delved into more technical detail. Substantial time was reserved for individual networking and communications.

Speakers were chosen based on their international recognition and covered various topics:

- **Ed Larsen** (University of Michigan) — Numerical methods used in neutron transport.
- **Tony Mezzacappa** (ORNL) — Applications and numerical methods used in supernova core collapse.

- **David Levermore** (University of Maryland) — Moment and closure approximations used in approximating transport equations.
- **Marty Marinak** (LLNL) — Transport needs in high-energy density physics.
- **George Kattawar** (Texas A&M) — Polarization and radiative transfer in oceanography.

The conference structure and venue work extremely well. Participants from different fields, who would never have had the opportunity to speak with other participants, were engaged in stimulating and very fruitful discussions. The atmosphere was collegial with all participants willing to learn and teach. As the conference week progressed, atmospheric scientists were learning about methods used in nuclear engineering.

A radiation physicist related, "Using the Fokker–Planck equation for studying our scattering problems never occurred to me. We are going to look into this." A mathematician who gave a talk on medical imaging and radiation oncology forged a bond with a nuclear engineer and is taking a sabbatical leave to apply his knowledge to radiation oncology. Another outgrowth of the workshop is that Ryan Clement (LLNL) and organizer Frank Graziani are setting up an e-print server that will serve as a repository for computational transport papers.

Overall, feedback from the workshop has been very positive and most expressed hope that it would be done again. The plan is to do a smaller, multi-disciplinary special-topics meeting next year followed by a larger meeting in 2006 similar to the 2004 workshop. Like the 2004 workshop, the future ones will be organized in conjunction with the Institute for Pure and Applied Mathematics at UCLA.

<http://www.ipam.ucla.edu/programs/tr2004/>

AMG / FOSLS Summit

Dates

September 27 – October 3, 2004

Location

Lake City, Colorado

In 1997, an informal meeting was held in Frisco, CO, between researchers in CU Boulder's Applied Mathematics Department and LLNL's CASC Division to discuss their collaboration on Algebraic Multigrid (AMG) methods. They met again in Boulder in 1998. In 2000, the meetings became annual, held in Lake City, CO. In 2002, they were expanded to include discussions on the First Order Systems Least Squares (FOSLS) methodology.

The summit is structured as a "working meeting" with an emphasis on exposing open research issues and generating ideas for solving them. Formal talks are strongly discouraged in favor of whiteboard discussions and individual interactions, a format that distinguishes it from typical meetings and conferences. Participation is by invitation only, consisting primarily of researchers from CU Boulder and CASC, but also including a small number of leading experts from other institutions around the world.

The Summit was held September 27–October 3, 2004. The first half of the meeting focused on AMG, and the second half focused on FOSLS, with an overlap day in between. There were 31 attendees this year. The main CU/CASC group consisted of 11 from CU Boulder, 9 from CASC, plus 4 recent CU graduates. The other attendees were: Irad Yavneh (Technion, Israel), Ludmil Zikatanov (Penn State), Ira Livshits (Ball State), Achi Brandt Weizmann Institute, Israel), Bobby Philip (LANL), Marzio Sala and Michael Gee (SNL).

The first half hour of the meeting was spent setting the agenda. The topics suggested this year by the attendees were as follows (as written on the whiteboard).

- Weighted Functionals
- Smooth Aggregation & e-Free AMGe
- Nonlinearity
- Almost Zero Modes
- Measures
- Coarse Variable Types
- Sparsity of P/Dilution
- Relaxation
- Sharp Theory
- Trace Minimization
- Iradism
- Upscaling
- Wavelet AMG

During the remainder of the meeting, each topic item (and its associated issues) was discussed in detail and solution approaches were proposed and debated. Two sample outcomes of the meeting were an improved theoretical foundation for a new trace minimization approach for defining interpolation in AMG, and an idea for relating a new compatible relaxation method (one that defines the coarse variables as averages) to a recent AMG theory and framework.